

**Amendments to the claims,  
Listing of all claims pursuant to 37 CFR 1.121(c)**

*This listing of claims will replace all prior versions, and listings, of claims in the application:*

What is claimed is:

1. (Currently amended) A method for capturing database changes at a primary database and applying those changes to a replicate database while the replicate database remains on-line and available for use, the method comprising:

monitoring transactions occurring at the primary database for detecting changes made to the primary database;

recording information about transactions observed to have occurred at the primary database in a transaction log;

synchronously copying the information about the transactions recorded in the transaction log to ~~the~~ a remote site, so as to create at the remote site a mirrored transaction log that is guaranteed to contain at a synchronized point in time an exact copy of the transactions recorded in the transaction log at the primary database;

while the replicate database remains on-line and available for use, replicating changes made at the primary database to the replicate database by:

reconstructing said transactions at the replicate database based on the information about the transactions copied to the mirrored transaction log; and

asynchronously applying the reconstructed transactions at the replicate database.

2. (Previously presented) The method of claim 1, wherein said transactions include selected ones of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.

3. (Previously presented) The method of claim 1, wherein said recording step includes recording at least one log record about one of the transactions in the transaction log.

4. (Original) The method of claim 3, wherein said at least one log record characterizes changes made to the primary database in the transaction.
5. (Original) The method of claim 1, wherein said synchronously copying step includes using a file mirroring module.
6. (Original) The method of claim 1, wherein said synchronously copying step includes using file replication hardware.
7. (Original) The method of claim 1, wherein said synchronously copying step includes using file replication software.
8. (Original) The method of claim 1, wherein said synchronously copying step includes synchronously copying information to the transaction log and the mirrored transaction log before completing the transaction at the primary database.
9. (Previously presented) The method of claim 1, wherein said synchronously copying step includes replicating at a file block level the information about the transactions in the transaction log to the mirrored transaction log.
10. (Original) The method of claim 1, further comprising:  
copying database schema information from the primary database to a site at which the mirrored transaction log is located to enable transactions to be reconstructed and applied at the replicate database.
11. (Previously presented) The method of claim 10, wherein said replicating step includes reconstructing said transactions at the replicate database based, at least in part, on said database schema information.
12. (Previously presented) The method of claim 1, wherein said replicating step includes formatting the reconstructed transactions to have the same format as the

transactions at the primary database.

13. (Previously presented) The method of claim 1, wherein said asynchronously applying step includes verifying that the reconstructed transactions are ordered correctly.

14. (Previously presented) The method of claim 1, wherein said asynchronously applying step includes applying the reconstructed transactions at the replicate database in the same order as the transaction order at the primary database.

15. (Original) The method of claim 1, further comprising:  
responding to a database query at the replicate database while a transaction is being replicated from the primary database to the replicate database.

16. (Original) A computer-readable medium having computer-executable instructions for performing the method of claim 1.

17. (Previously presented) The method of claim 1, further comprising:  
downloading a set of computer-executable instructions for performing the method of claim 1.

18. (Previously presented) A system for replicating transactions from a source database to a standby database, the system comprising:  
source and replicate databases sharing a common schema, the source database having a transaction log for recording log records for transactions performed at the source database;  
a mirrored transaction log for recording mirror copies of the log records for transactions performed at the source database, so as to create at a remote site an exact synchronous copy of the transaction log at a given point in time;  
a file mirroring module for synchronously replicating log records from the transaction log to the mirrored transaction log as transactions are performed at the source database;

a log reader module at the remote site for reading log records in the mirrored transaction log and reconstructing transactions for application to the standby database based upon log records in the mirrored transaction log; and

a distribution module for asynchronously applying the transactions reconstructed by the log reader module to the standby database;

wherein both the source and standby databases remain on-line for use while transactions are replicated from one to the other.

19. (Original) The system of claim 18, wherein said standby database is available for responding to database queries while transactions are being replicated from the source database to the standby database.

20. (Original) The system of claim 18, wherein said transactions include a selected one of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.

21. (Original) The system of claim 18, wherein said log records characterize changes made to the source database based upon transactions performed at the source database.

22. (Original) The system of claim 18, wherein said file mirroring module comprises file replication hardware.

23. (Original) The system of claim 18, wherein said file mirroring module comprises a disk mirroring module.

24. (Original) The system of claim 18, wherein said file mirroring module replicates log records in the transaction log to the mirrored transaction log at a file block level.

25. (Original) The system of claim 18, wherein said file mirroring module

replicates log records relating to a particular transaction performed at the source database to the mirrored transaction log before said particular transaction is completed at the source database.

26. (Original) The system of claim 18, wherein said log reader module reconstructs transactions based, at least in part, on database schema information for the source database.

27. (Previously presented) The system of claim 26, further comprising: database schema information for the source database, which is copied to the remote site to enable transactions to be reconstructed and applied at the standby database.

28. (Original) The system of claim 18, wherein said log reader module formats the reconstructed transactions so that the reconstructed transactions are in the same format as the transaction at the source database.

29. (Original) The system of claim 18, wherein said distribution module applies reconstructed transactions at the standby database in the same order as the order of transactions applied at the source database.

30. (Previously presented) A method for replicating database operations from a first database to a second database while making the second database available for decision support purposes, the method comprising:

as database operations are performed at the first database, generating a first set of log records characterizing said operations and recording them in a first log, so that the first log reflects all changes made to the first database;

as the first set of log records are recorded in the first log, synchronously recording in a second log located at remote site a second set of log records that comprise an exact copy of the first set of log records, so that said second log comprises at a synchronized point in time an exact copy of said first log; and

while the second database is available for decision support purposes,

asynchronously replicating said operations performed at the first database at the second database by asynchronously performing the substeps of:

constructing replicate operations based, at least in part, on said second set of log records recorded in the second log; and  
applying the replicate operations at the second database.

31. (Previously presented) The method of claim 30, wherein said operations include selected ones of a Structured Query Language (SQL) "INSERT", "UPDATE", "DELETE", "DDL" and "Procedure" operation.

32. (Original) The method of claim 30, wherein said synchronously recording step includes file mirroring.

33. (Original) The method of claim 30, wherein said synchronously recording step includes replicating said at least one log record to the second log at a file block level.

34. (Original) The method of claim 30, wherein said synchronously recording step includes using a disk mirroring module.

35. (Original) The method of claim 30, further comprising:  
copying database schema information from the first database prior to performing said operation at the first database.

36. (Original) The method of claim 35, wherein said constructing substep includes constructing a replicate operation based, at least in part, on said database schema information.

37. (Original) The method of claim 35, further comprising:  
tracking modifications to said database schema information at the first database;  
and  
constructing a replicate operation based on said database schema information in

effect when the operation is performed at the first database.

38. (Original) The method of claim 30, further comprising:  
assigning a unique identifier to database objects at the first database;  
if a database object is modified, assigning a different unique identifier to the  
database object that is modified; and  
determining a particular database object to be used in constructing a replicate  
operation based upon said unique identifier assigned to said particular database object.

39. (Previously presented) The method of claim 30, wherein said constructing  
substep includes formatting the replicate operations in the same manner as said  
operations at the first database.

40. (Previously presented) The method of claim 30, wherein said applying  
substep includes applying the replicate operations at the second database in the same  
order as said operations are applied at the first database.

41. (Original) The method of claim 30, wherein making the second database  
available for decision support purposes includes responding to a database query as said  
operation is being replicated.

42. (Original) The method of claim 30, wherein making the second database  
available for decision support purposes includes providing access to data in the second  
database as said operation is being replicated.

43. (Previously presented) A method for replicating transactions from a primary  
database to a replicate database while the replicate database remains available for use, the  
primary database and the replicate database sharing a common schema, the method  
comprising:

recording log records for transactions being performed at a primary database in a  
primary transaction log;

synchronously creating a mirrored transaction log at a remote site, the mirrored transaction log comprising an exact and synchronized copy of the log records in the primary transaction log;

generating reconstructed transactions based on the copies of the log records in the mirrored transaction log at the remote site; and

asynchronously applying the reconstructed transactions at the replicate database while the replicate database remains available for use.